

‘local’ vs. ‘global’ parameters – breaking the gaussian complexity barrier

Shahar Mendelson*

April 10, 2015

Abstract

We show that if F is a convex class of functions that is L -subgaussian, the error rate of learning problems generated by independent noise is equivalent to a fixed point determined by ‘local’ covering estimates of the class, rather than by the gaussian averages. To that end, we establish new sharp upper and lower estimates on the error rate for such problems.

1 Introduction

The focus of this article is on the question of *prediction*. Given a class of functions F defined on a probability space (Ω, μ) and an unknown target random variable Y , one would like to identify an element of F whose ‘predictive capabilities’ are (almost) the best possible in the class. The notion of ‘best’ is measured via the point-wise cost of predicting $f(x)$ instead of y , and the best function in the class is the one that minimizes the average cost. Here, we will consider the squared loss: the cost of predicting $f(x)$ rather than y is $(f(x) - y)^2$, and if X is distributed according to μ , the goal is to identify

$$f^* = \operatorname{argmin}_{f \in F} \mathbb{E}(f(X) - Y)^2 = \operatorname{argmin}_{f \in F} \|f - Y\|_{L_2}^2,$$

where the expectation is taken with respect to the joint distribution of X and Y on the product space $\Omega \times \mathbb{R}$.

*Department of Mathematics, Technion, I.I.T, Haifa 32000, Israel
email: shahar@tx.technion.ac.il

Supported in part by the Mathematical Sciences Institute, The Australian National University, Canberra, ACT 2601, Australia. Additional support was given by the Israel Science Foundation grant 900/10.

The information at one's disposal is rather limited: a random sample $(X_i, Y_i)_{i=1}^N$, selected according to the N -product of the joint distribution of X and Y . And, using this data, one must select some (random) $f \in F$.

Definition 1.1 *Given a sample size N and a class F defined on (Ω, μ) , a learning procedure is a map $\Psi : (\Omega \times \mathbb{R})^N \rightarrow F$. For a set \mathcal{Y} of admissible targets, Ψ performs with confidence $1-\delta$ and accuracy \mathcal{E}_p if for every $Y \in \mathcal{Y}$, and setting $\tilde{f} = \Psi((X_i, Y_i)_{i=1}^N)$,*

$$\mathbb{E} \left((\tilde{f} - Y)^2 | (X_i, Y_i)_{i=1}^N \right) \leq \mathbb{E}(f^*(X) - Y)^2 + \mathcal{E}_p$$

with probability at least $1-\delta$ relative to the N -product of the joint distribution of X and Y .

The accuracy (or error) \mathcal{E}_p is a function of F , N and δ , and may depend on some features of the target Y as well, for example, its norm in some L_q space.

A fundamental problem in Learning Theory is to identify the features of the underlying class F and of the set of admissible targets \mathcal{Y} that govern \mathcal{E}_p ; in particular, the way \mathcal{E}_p scales with the sample size N (the so-called *error rate*). This question has been studied extensively, and we refer the reader to the manuscripts [2, 9, 6, 17, 3, 3, 10, 11] for more information on its history and on some more recent progress.

Here, the aim is to obtain matching upper and lower bounds on \mathcal{E}_p that hold for any reasonable class F , at least under some assumptions which we will now outline.

It is well understood that the ability to predict is quantified by various complexity parameters of the underlying class. Frequently, one encounters parameters that are based on various gaussian and empirical/multiplier processes indexed by ‘localizations’ of F (see, e.g., [7]), and any hope of obtaining matching bounds on \mathcal{E}_p must be based on sharp estimates on these processes. Unfortunately, the analysis of empirical/multiplier processes is, in general, highly nontrivial. Moreover, and unlike gaussian processes, there is no clear path that leads to sharp bounds on empirical processes, and even when upper estimates are available, they are often loose and lead to suboptimal bounds on \mathcal{E}_p .

The one generic example in which a more satisfactory theory of empirical/multiplier processes is known, is when the indexing class is L -subgaussian.

Definition 1.2 A class $F \subset L_2(\mu)$ is L -subgaussian with respect to the measure μ if for every $p \geq 2$ and every $f, h \in F \cup \{0\}$,

$$\|f - h\|_{L_p(\mu)} \leq L\sqrt{p}\|f - h\|_{L_2(\mu)},$$

and if the canonical gaussian process $\{G_f : f \in F\}$ is bounded (see the book [4] for a detailed survey on gaussian processes).

More facts on subgaussian classes may be found in [8, 18, 4, 14, 7]. For our purposes, the main feature of subgaussian classes is that the empirical and multiplier processes that govern \mathcal{E}_p may be bounded from above using properties of the canonical gaussian process indexed by the class, giving one some hope of obtaining sharp estimates. Because of that feature, we will focus in what follows on subgaussian classes.

Despite their importance, complexity parameters are not the entire story when it comes to \mathcal{E}_p . For example, it is possible to construct a class consisting of just two functions, $\{f_1, f_2\}$, but if the target Y is a $1/\sqrt{N}$ -perturbation of the midpoint $(f_1 + f_2)/2$, no learning procedure can perform with an error that is better than c/\sqrt{N} having been given a sample of cardinality N (see, e.g., [1]). Thus, rather than being solely determined by the complexity of the underlying class, there is an additional geometric requirement on F and \mathcal{Y} which is there to ensure that all the admissible targets in \mathcal{Y} are located in a favourable position relative of F (see [13] for more details). One may show that if $F \subset L_2(\mu)$ is compact and convex, any target $Y \in L_2$ is in a favourable position relative to F . Therefore, to remove possible geometric obstructions, we will assume that $F \subset L_2(\mu)$ is compact and convex.

Finally, for a reason that will become clear later, we will not study a general class of admissible targets \mathcal{Y} , but rather consider targets of the form $Y = f(X) + W$ for some $f \in F$ and W that is orthogonal to $\text{span}(F)$ (e.g., $W \in L_2$ that is a mean-zero random variable and is independent of X is a ‘legal’ choice).

With all these assumptions in place, let us formulate the question we would like to study:

Question 1.3 Let $F \subset L_2(\mu)$ be a compact, convex class that is L -subgaussian with respect to μ . Given targets of the form $Y = f(X) + W$ as above, find matching upper and lower bounds (up to constants) on \mathcal{E}_p .

Let us recall the following standard definitions.

Definition 1.4 Let $F \subset L_2(\mu)$. Set

$$F - h = \{f - h : f \in F\} \quad \text{and} \quad F - F = \{f - h : f, h \in F\}.$$

Denote by

$$\text{star}(F) = \{\lambda f : f \in F, 0 \leq \lambda \leq 1\}$$

the star-shaped hull of F with 0; F is star-shaped around 0 if $\text{star}(F) = F$.

Let $\{G_f : f \in F\}$ be the canonical gaussian process indexed by F and set

$$\mathbb{E}\|G\|_F = \sup \left\{ \mathbb{E} \sup_{f \in F'} G_f : F' \subset F, F' \text{ is finite} \right\}.$$

Finally, let D be the unit ball in $L_2(\mu)$.

The best known bounds on \mathcal{E}_p in the subgaussian context have been established in [7] and are based on two fixed points:

Definition 1.5 For $\kappa_1, \kappa_2 > 0$, set

$$r_M(\kappa_1, f) = \inf \left\{ s > 0 : \mathbb{E}\|G\|_{(F-f) \cap sD} \leq \kappa_1 s^2 \sqrt{N} \right\} \quad (1.1)$$

and

$$r_Q(\kappa_2, f) = \inf \left\{ r > 0 : \mathbb{E}\|G\|_{(F-f) \cap sD} \leq \kappa_2 s \sqrt{N} \right\}. \quad (1.2)$$

Put

$$r_M(\kappa_1) = \sup_{f \in F} r_M(\kappa_1, f) \quad \text{and} \quad r_Q(\kappa_2) = \sup_{f \in F} r_Q(\kappa_2, f).$$

In the context of the problem we are interested in, one has the following:

Theorem 1.6 [7] For every $L \geq 1$ there exist constants c_1, c_2, c_3 and c_4 that depend only on L for which the following holds. Let $F \subset L_2(\mu)$ be a compact, convex, L -subgaussian class of functions, set $Y = f_0(X) + W$ and assume that for every $p \geq 2$, $\|W\|_{L_p} \leq L\sqrt{p}\|W\|_{L_2}$. There is a learning procedure (empirical risk minimization performed in F) for which, if

$$r \geq 2 \max \{r_M(c_0/\|W\|_{L_2}), r_Q(c_1)\}.$$

then with probability at least

$$1 - 2 \exp \left(-c_2 N \min \{1, r^2 / \|W\|_{L_2}^2\} \right),$$

the error of the procedure is at most $\mathcal{E}_p \leq r^2$.

The lower bound that complements Theorem 1.6 uses ‘local’ analogs of r_M and r_Q that are based on the notion of packing numbers.

Definition 1.7 *Let E be a normed space and set B to be its unit ball. Let $\mathcal{M}(A, rB)$ be the cardinality of a maximal r -separated subset of A with respect to the given norm, that is, the cardinality of the largest subset $(a_i)_{i=1}^m \subset A$ for which $\|a_i - a_j\| \geq r$ for every $i \neq j$.*

Definition 1.8 *For $\eta_1, \eta_2 > 0$ set*

$$\gamma_M(\eta_1, f) = \inf \{s > 0 : \log \mathcal{M}((F - f) \cap 4sD, (s/2)D) \leq \eta_1^2 s^2 N\}.$$

and

$$\gamma_Q(\eta_2, f) = \inf \{s > 0 : \log \mathcal{M}((F - f) \cap 4sD, (s/2)D) \leq \eta_2^2 N\}.$$

Put

$$\gamma_M(\eta_1) = \sup_{f \in F} \gamma_M(\eta_1, f), \quad \text{and} \quad \gamma_Q(\eta_2) = \sup_{f \in F} \gamma_Q(\eta_2, f).$$

Theorem 1.9 [7] *There exist absolute constants c_1 and c_2 for which the following holds. Let F be a class of functions, set W be a centred normal random variable and for every $f \in F$ put $Y^f = f(X) + W$. If Ψ is a learning procedure that performs for every target Y^f with confidence at least $3/4$, then there is some Y^f for which $\mathcal{E}_p \geq c_1 \gamma_M^2(c_2/\|W\|_{L_2})$.*

Remark 1.10 *One should note that a lower bound that is based on γ_Q was not known.*

The connection between the two types of parameters is Sudakov’s inequality (see, e.g. [8]): there is an absolute constant c for which, for every $H \subset L_2(\mu)$,

$$c \sup_{\varepsilon > 0} \varepsilon \log^{1/2} \mathcal{M}(H, \varepsilon D) \leq \mathbb{E}\|G\|_H.$$

To see the connection, assume that for every $f \in F$, $\mathbb{E}\|G\|_{(F-f) \cap 4rD} \leq \kappa_1(4r)^2 \sqrt{N}$, which means that $r_M(\kappa_1) \leq 4r$. Applying Sudakov’s inequality to $H = (F - f) \cap 4rD$ and for the choice of $\varepsilon = r/2$,

$$c(r/2) \log^{1/2} \mathcal{M}((F - f) \cap 4rD, rD) \leq \mathbb{E}\|G\|_{(F-f) \cap 4rD} \leq 16\kappa_1 r^2 \sqrt{N};$$

hence, $\gamma_M(c_1 \kappa_1) \leq r$. A similar observation is true for r_Q and γ_Q , which shows that γ_M and γ_Q are intrinsically smaller than r_Q and r_M respectively, for the right choice of constants.

The starting point of this article is fact that the gap between these upper and lower estimates on \mathcal{E}_p is more than a mere technicality.

The core issue is that the parameters r_M and r_Q are ‘global’ in nature, whereas γ_M and γ_Q are ‘local’. Indeed, although $(F - f) \cap rD$ is a localized set, $\mathbb{E}\|G\|_{(F-f) \cap rD}$ is not determined solely by the effects of a ‘level’ that is proportional r . For example, it is straightforward to construct examples in which $\mathbb{E}\|G\|_{(F-f) \cap rD} \geq cr\sqrt{N}$ because of a very large, ρ -separated subset of $(F - f) \cap rD$, for ρ that is much smaller than r . Thus, even if r_M or r_Q are of order r , this need not be ‘exhibited’ by $(F - f) \cap rD$ at a scale that is proportional to r . In contrast, γ_M and γ_Q are ‘local’: the degree of separation is proportional to the diameter of the separated set, and the fixed point indicates that $(F - f) \cap rD$ is truly ‘rich’ at a scale that is proportional to r .

As noted in [7], the upper and lower estimates coincide when the ‘local’ and ‘global’ parameters are equivalent, but that is not a typical situation – in the generic case, there is a gap between the two. An example of that fact will be presented in Section 5.

Given that there is a gap between the two sets of parameters, one must face the obvious question: which of the two captures \mathcal{E}_p ? Is it the ‘global’ pair, r_Q and r_M , or the ‘local’ one of γ_Q and γ_M ?

Our main result is that the ‘local’ parameters are the right answer – at least in the setup outlined above. To that end, we shall improve the upper bound in Theorem 1.6 and add the missing component in Theorem 1.9.

Theorem 1.11 *For every $L > 1$ and $q > 2$ there are constants c_0, \dots, c_5 that depend only of q and L for which the following holds. Let $F \subset L_2(\mu)$ be a compact, convex, L -subgaussian class of functions with respect to μ . There is a learning procedure $\Psi : (\Omega \times \mathbb{R})^N \rightarrow F$, for which, if $Y = f(X) + W$ for $f \in F$ and $W \in L_q$ that is orthogonal to $\text{span}(F)$, then with probability at least*

$$1 - 2 \exp \left(-c_0 N \min\{1, \gamma_M^2(c_1/\|W\|_{L_q})\} \right) - c_2 \frac{\log^q N}{N^{(q/2)-1}},$$

$$\mathcal{E}_p \leq c_3 \max \left\{ \gamma_M^2 \left(\frac{c_1}{\|W\|_{L_q}} \right), \gamma_Q^2(c_4) \right\} + r_Q^2(c_4) \exp(-c_5 \exp(N))$$

The term $r_Q^2(c_4) \exp(-c_5 \exp(N))$ is almost certainly an artifact of the proof, but in any case, it is significantly smaller than the dominating term in any reasonable example.

To complement Theorem 1.11 we obtain the following lower bound.

Theorem 1.12 *There exist absolute constants c_0 and c_1 for which the following holds. Let $F \subset L_2(\mu)$ be a convex, centrally-symmetric class of functions and let Ψ be any learning procedure that performs with confidence $7/8$ for any target of the form $Y = f(X) + W$ for some $f \in F$ and $W \in L_2$ that is orthogonal to $\text{span}(F)$.*

- *For any $W \in L_2$ that is orthogonal to $\text{span}(F)$, there is some $f \in F$, for which, for $Y = f(X) + W$,*

$$\mathcal{E}_p \geq c_0 \gamma_Q^2(c_1).$$

- *If W is a centred, normal random variable that is independent of X , there is some $f \in F$ for which, for $Y = f(X) + W$,*

$$\mathcal{E}_p \geq c_0 \gamma_M^2 \left(\frac{c_1}{\|W\|_{L_2}} \right).$$

An outcome of Theorem 1.11 and Theorem 1.12 is that if W is a centred gaussian random variable that is independent of X , then for any convex, centrally-symmetric, L -subgaussian class F , the upper and lower estimates match (up to the parasitic and negligible term $r_Q^2(c_4) \exp(-c_5 \exp(N))$ in the upper bound): when considering targets of the form $Y = f(X) + W$ for $f \in F$,

$$\mathcal{E}_p \sim \max \{ \gamma_Q^2(c_1), \gamma_M^2(c_2/\|W\|_{L_2}) \}.$$

The second part of Theorem 1.12 follows from Theorem 1.9. We have chosen to present a new proof of that fact – a proof we believe is both instructive and less restrictive than existing proofs. The first part of Theorem 1.12 is, to the best of our knowledge, new.

Let us mention that if F happens to be convex and centrally symmetric (i.e. if $f \in F$ then $-f \in F$), what is essentially the ‘richest’ shift of F is the 0-shift. Indeed, since $F - F = 2F$, it is evident that for every $f \in F$

$$(F - f) \cap 4rD \subset (F - F) \cap 4rD = 2(F \cap 2rD).$$

This makes one’s life much simpler when studying lower bounds, as it gives an obvious choice of where to look. Indeed, the ‘richest’ part of F is the hardest part for a learning procedure to deal with – and that part is a neighbourhood of 0.

1.1 The idea of the proof of the upper bound

The proof of the upper bound is based on the following decomposition of the squared excess loss: let Y be the unknown target and set $f^* = \operatorname{argmin}_{f \in F} \|f - Y\|_{L_2}$. For every $f \in F$, let $\ell_f(X, Y) = (f(X) - Y)^2$ and set

$$\begin{aligned}\mathcal{L}_f^F(X, Y) &= (\ell_f - \ell_{f^*})(X, Y) = (f(X) - Y)^2 - (f^*(X) - Y)^2 \\ &= 2(f^*(X) - Y)(f - f^*)(X) + (f - f^*)^2(X).\end{aligned}\quad (1.3)$$

Let $P_N h = \frac{1}{N} \sum_{i=1}^N h(X_i, Y_i)$ and set

$$\hat{f} = \operatorname{argmin}_{f \in F} P_N \ell_f = \operatorname{argmin}_{f \in F} P_N \mathcal{L}_f^F$$

to be the empirical minimizer in F . The learning procedure that assigns to every sample $(X_i, Y_i)_{i=1}^N$ the empirical minimizer in F is called *Empirical Risk Minimization* (ERM).

Clearly, $\mathcal{L}_{f^*}^F = 0$, and thus, for every sample $(X_i, Y_i)_{i=1}^N$,

$$P_N \mathcal{L}_{\hat{f}}^F \leq 0,$$

implying that members of the random set $\{f \in F : P_N \mathcal{L}_f^F > 0\}$ cannot be empirical minimizers. One way of identifying that set is via the decomposition (1.3): assume that $(X_i, Y_i)_{i=1}^N$ is a sample for which, if $\|f - f^*\| \geq r$, one has

$$\frac{1}{N} \sum_{i=1}^N (f - f^*)^2 \geq \kappa \|f - f^*\|_{L_2}^2, \quad (1.4)$$

and

$$\left| \frac{1}{N} \sum_{i=1}^N (f^*(X_i) - Y_i)(f - f^*)(X_i) - \mathbb{E}(f^*(X) - Y)(f - f^*)(X) \right| \leq \frac{\kappa}{4} \|f - f^*\|_{L_2}^2. \quad (1.5)$$

Since F is compact and convex, by properties of the metric projection onto a closed convex set in an inner product space,

$$\mathbb{E}(f^*(X) - Y)(f - f^*)(X) \geq 0 \quad (1.6)$$

for every $f \in F$. Therefore, setting $\xi = f^*(X) - Y$ and $\xi_i = f^*(X_i) - Y_i$,

$$\begin{aligned}P_N \mathcal{L}_f^F &\geq \frac{1}{N} \sum_{i=1}^N (f - f^*)^2 - 2 \left| \frac{1}{N} \sum_{i=1}^N \xi_i (f - f^*)(X_i) - \mathbb{E} \xi (f - f^*)(X) \right| \\ &\quad + \mathbb{E} \xi (f - f^*)(X) \geq \kappa - 2(\kappa/4) > 0\end{aligned}$$

for every $f \in F$ that satisfies $\|f - f^*\|_{L_2} \geq r$. Thus, if (1.4) and (1.5) hold for the sample $(X_i, Y_i)_{i=1}^N$, then

$$\{f \in F : \|f - f^*\|_{L_2} \geq r\} \subset \{f \in F : P_N \mathcal{L}_f^F > 0\}$$

implying that $\|\hat{f} - f^*\|_{L_2} < r$.

This argument has been used in [10] and was then extended in [11], showing that

$$\mathbb{E}(\mathcal{L}_{\hat{f}}^F | (X_i, Y_i)_{i=1}^N) \leq r^2$$

– which is the type of result one is looking for.

This method of proof leads to the complexity parameters r_Q and r_M : the former controls the quadratic component (1.4) and the latter the multiplier component (1.5). The ‘global’ nature of r_Q and r_M , i.e., the fact that the two depend on the gaussian oscillation $\mathbb{E}\|G\|_{(F-f) \cap rD}$ cannot be helped: the oscillations of the quadratic and multiplier processes are highly affected by the ‘richness’ of F around f^* at every ‘level’.

A rather obvious idea for improving the upper estimate is ‘erasing’ all the fine structure of F , for example, by replacing F with an appropriate separated subset. The difficulty in such an approach is that the geometry of a separated set is problematic, and (1.6) will no longer be true for an arbitrary target Y . This is why we only consider targets of the form $f(X) + W$ for $f \in F$ and W that is orthogonal to $\text{span}(F)$. For such targets, a version of (1.6) happens to be true even if F is replaced by a separated set.

The path we will take in proving the upper bound is as follows:

- Choose a ‘correct’ level r using the parameters γ_M and γ_Q for well-chosen constants η_1 and η_2 that depend only on q and L .
- Replace F by V , a maximal r -separated subset of F with respect to the $L_2(\mu)$ norm, and study ERM in V . To that end, set $v_0 = \operatorname{argmin}_{v \in V} \|v - Y\|_{L_2}$ and observe that by the orthogonality of W to $\text{span}(F)$, for every $v \in V$,

$$|\mathbb{E}(v_0(X) - Y)(v - v_0)(X)| = |\mathbb{E}(v_0 - f^*)(v - v_0)(X)| \leq r\|v - v_0\|_{L_2}.$$

Therefore, the empirical excess loss relative to V satisfies

$$\begin{aligned} P_N \mathcal{L}_v^V &\geq \frac{1}{N} \sum_{i=1}^N (v - v_0)^2(X_i) \\ &\quad - 2 \left| \frac{1}{N} \sum_{i=1}^N (v_0(X_i) - Y_i)(v - v_0)(X_i) - \mathbb{E}(v_0(X) - Y)(v - v_0)(X) \right| \\ &\quad - 2r \|v - v_0\|_{L_2}. \end{aligned}$$

- Next, one may study the corresponding quadratic and multiplier processes indexed by localizations of V and show that with high probability, if $\|v - v_0\|_{L_2} \geq c_1 r$ then $P_N \mathcal{L}_v^V > 0$. Thus, ERM performed in V produces \hat{v} for which $\|\hat{v} - v_0\|_{L_2} \leq c_1 r$.
- It is possible to show that on the same event, $\|\hat{v} - f^*\|_{L_2} \leq c_2 r$. And, using the orthogonality of W to $\text{span}(F)$ once again, $\mathbb{E}(\mathcal{L}_{\hat{v}}^F | (X_i, Y_i)_{i=1}^N) \leq c_3 r^2$, as required.

2 Preliminaries

Let us begin with some notation. Throughout, absolute constants are denoted by c, c_1, \dots etc. Their value may change from line to line. $c(\alpha)$ is a constant that depends only on the parameter α . We use $\kappa_1, \kappa_2, \eta_1, \eta_2$ etc. to denote fixed constants whose value remains unchanged throughout the article.

In what follows, we will, at times, abuse notation and not specify the probability space on which each random variable is defined. For example, $\|f - Y\|_{L_2}^2 = \mathbb{E}(f(X) - Y)^2$ and integration is with respect to the joint distribution of X and Y , while $\|f - f_0\|_{L_2}^2 = \mathbb{E}(f - f_0)^2(X)$, in which case integration is with respect to μ .

Next, let us turn to the notions of *cover* and *covering numbers*.

Definition 2.1 *Let B be a unit ball of a norm. Set $\mathcal{N}(A, B)$ to be the minimal number of centres $a_1, \dots, a_n \in A$ for which $A \subset \bigcup_{i=1}^n (a_i + B)$. $(a_i)_{i=1}^n$ is called a cover of A with respect to B . An r -cover is a cover with respect to the set rB .*

It is standard to verify that if a_1, \dots, a_m is a maximal separated subset with respect to B then it is also a cover with respect to B . Indeed, the maximality of the separated set implies that every point $a \in A$ has some a_i for which

$\|a_i - a\| \leq 1$, i.e. $a \in a_i + B$. Therefore, $\mathcal{N}(A, B) \leq \mathcal{M}(A, B)$. In the reverse direction, if a_1, \dots, a_n is a cover with respect to B , then each one of the balls $a_i + B$ contains at most one point in any 2-separated set. Thus, $\mathcal{M}(A, 2B) \leq \mathcal{N}(A, B)$.

The following lemma is straightforward but it plays a crucial part in what follows.

Lemma 2.2 *Let $T \subset W \subset L_2(\mu)$. For $s > r > 0$, set*

$$\phi(s, r) = \sup_{w \in W} \mathcal{N}(T \cap (w + sD), rD).$$

Then

1. $\phi(s, r) \leq \phi(s, s/2) \cdot \phi(s/2, r)$.
2. *If T and W are star-shaped around 0 then*

$$\log \phi(s, r) \leq c_0 \log(2s/r) \cdot \log \phi(4r, r)$$

for a suitable absolute constant c_0 .

Proof. Fix $w \in W$ and let $t_1, \dots, t_N \in T \cap (w + sD)$ be centres of a minimal $s/2$ -cover of that set. For every $1 \leq i \leq N$,

$$T \cap (w + sD) \cap (t_i + (s/2)D) \subset T \cap (t_i + (s/2)D),$$

and $\mathcal{N}(T \cap (t_i + (s/2)D), rD) \leq \phi(s/2, r)$, because $t_i \in T \subset W$. Therefore,

$$\sup_{w \in W} \mathcal{N}(T \cap (w + sD), rD) \leq \sup_{w \in W} \mathcal{N}(T \cap (w + sD), (s/2)D) \cdot \phi(s/2, r).$$

Turning to the second part of the claim, assume that T and W are star-shaped around 0. Let $w \in W$, set t_1, \dots, t_m to be a maximal $s/2$ -separated subset of $T \cap (w + sD)$ with respect to the $L_2(\mu)$ norm and put $y_i = (r/s)t_i$. Since T is star-shaped around 0, $y_i \in T$ and $(y_i)_{i=1}^m$ is an $r/2$ -separated subset of $(r/s)w + rD$. For the same reason, $(r/s)w \in W$, and

$$\mathcal{M}(T \cap (w + sD), rD) \leq \sup_{v \in W} \mathcal{M}(T \cap (v + rD), (r/2)D).$$

Using the standard connection between packing numbers and covering numbers and taking the supremum over w ,

$$\begin{aligned} \phi(s, s/2) &= \sup_{w \in W} \mathcal{N}(T \cap (w + sD), (s/2)D) \leq \sup_{w \in W} \mathcal{M}(T \cap (w + sD), (s/2)D) \\ &\leq \sup_{w \in W} \mathcal{M}(T \cap (w + 2rD), rD). \end{aligned}$$

Iterating the first part of the lemma,

$$\begin{aligned}
\log \phi(s, r) &\leq \log_2(2s/r) \cdot \sup_{w \in W} \log \mathcal{M}(T \cap (w + 4rD), 2rD) \\
&\leq \log_2(2s/r) \cdot \sup_{w \in W} \log \mathcal{N}(T \cap (w + 4rD), rD) \\
&\leq \log_2(2s/r) \cdot \log \phi(4r, r),
\end{aligned}$$

as claimed. ■

Before we turn to the proof of the upper bound, let us revisit the complexity parameters in question. Since F is a convex class, $F - f$ is star-shaped around 0; hence, if $s > r$

$$\mathcal{M}((F - f) \cap 4sD, (s/2)D) \leq \mathcal{M}((F - f) \cap 4(r/2)D, rD).$$

In particular, if $\gamma_M(\eta_1, f) < r$ then

$$\log \mathcal{M}((F - f) \cap 4sD, (s/2)D) \leq \eta_1^2 N r^2 \leq \eta_1^2 N s^2,$$

implying that $\gamma_M(\eta_1, f) < s$ as well.

This simple argument shows that if $r < \gamma_M(\eta_1, f)$ then

$$\log \mathcal{M}((F - f) \cap 4rD, (r/2)D) \geq \eta_1^2 N r^2,$$

while if $r > \gamma_M(\eta_1, f)$, the reverse inequality holds.

A similar assertion holds for γ_Q , r_M and r_Q ; the rather standard proof of these facts, which is almost identical to the argument used above, is omitted.

3 The upper bound

Let $F \subset L_2(\mu)$ be a compact, convex class of functions. Fix $r > 0$ that will be named later and let V to be a maximal r -separated subset of F . Note that for every $v_0 \in V$, $F_{v_0} = F - v_0$ is star-shaped around 0, and $\text{star}(V - v_0) \subset F - v_0$. Using the notation of Lemma 2.2, let $T = W = F_{v_0}$, and for $s > 2r > 0$,

$$\begin{aligned}
&\log \mathcal{N}((\text{star}(V - v_0)) \cap sD, rD) \leq \log \mathcal{N}(F_{v_0} \cap sD, rD) \\
&\leq \sup_{x \in F} \log \mathcal{N}(F_{v_0} \cap (x - v_0 + sD), rD) \\
&\leq c_0 \log(s/r) \sup_{x \in F} \log \mathcal{N}(F_{v_0} \cap (x - v_0 + 4rD), rD) \\
&= c_0 \log(s/r) \sup_{x \in F} \log \mathcal{N}(F \cap (x + 4rD), rD).
\end{aligned}$$

Also, observe that $F \cap (x + 4rD) \subset ((F - x) \cap 4rD) + x$, implying that

$$\log \mathcal{N}((\text{star}(V - v_0)) \cap sD, rD) \leq c_0 \log(s/r) \cdot \sup_{x \in F} \log \mathcal{N}((F - x) \cap 4rD, rD). \quad (3.1)$$

Moreover, the same estimate holds for $(V - v_0) \cap sD$, and since $V - v_0$ is r -separated,

$$\begin{aligned} \log |(V - v_0) \cap sD| &= \log \mathcal{M}((V - v_0) \cap sD, rD) \leq \log \mathcal{N}((V - v_0) \cap sD, (r/2)D) \\ &\leq \log \mathcal{N}(F_{v_0} \cap sD, (r/2)D) \\ &\leq c_0 \log(s/r) \cdot \sup_{x \in F} \log \mathcal{N}((F - x) \cap 4rD, (r/2)D) \\ &\leq c_0 \log(s/r) \cdot \sup_{x \in F} \log \mathcal{M}((F - x) \cap 4rD, (r/2)D) \end{aligned} \quad (3.2)$$

With that in mind, fix constants η_1, η_2, κ_2 and κ_3 that will be specified later, and for that choice of constants, let $r > 0$ for which

$$\sup_{x \in F} \log \mathcal{M}((F - x) \cap 4rD, (r/2)D) \leq \max \{ \eta_1^2 N r^2, \eta_2^2 N \}, \quad (3.3)$$

and

$$r \geq r_Q(\kappa_2) \exp(-\kappa_3 \exp(N));$$

that is,

$$r \geq \max \{ \gamma_M(\eta_1), \gamma_Q(\eta_2), r_Q(\kappa_2) \exp(-\kappa_3 \exp(N)) \}.$$

Let V be a maximal r -separated subset of F with respect to the $L_2(\mu)$ norm. Following the path outlined earlier, the idea is to study ERM in V , given the data $(X_i, Y_i)_{i=1}^N$ for $Y = f_0(X) + W$. To that end, one must control the multiplier and quadratic components in the decomposition of the squared loss relative to V : if $v_0 = \arg\min_{v \in V} \|v(X) - Y\|_{L_2}$,

$$\begin{aligned} \mathcal{L}_v^V(X, Y) &= (v(X) - Y)^2 - (v_0(X) - Y)^2 \\ &= 2(v_0(X) - Y)(v - v_0)(X) + (v - v_0)^2(X). \end{aligned}$$

Let us begin with the multiplier component:

Lemma 3.1 *Fix $0 < \theta < 1$, $L > 1$ and $q > 2$. There exist constants c_0, c_1 and c_2 that depend only on L and q and for which the following holds. Let F be a convex, L -subgaussian class, set $\xi \in L_q$ for some $q > 2$ and put $\eta_1 = c_0 \theta / \|\xi\|_{L_q}$. Then, for every $v_0 \in V$, with probability at least*

$$1 - c_1 \frac{\log^q N}{N^{((q/2)-1)}} - 2 \exp(-c_2 \eta_1^2 r^2 N),$$

$$\sup_{\{v \in V : \|v - v_0\|_{L_2} \geq 2r\}} \left| \frac{1}{N} \sum_{i=1}^N \xi_i \frac{v - v_0}{\|v - v_0\|_{L_2}^2} (X_i) - \mathbb{E} \xi \frac{v - v_0}{\|v - v_0\|_{L_2}^2} \right| \leq \theta.$$

The proof of Lemma 3.1 is based on the following fact from [12].

Theorem 3.2 *For $L > 1$ and $q > 2$ there exist constants c_0, c_1 and c_2 that depend only on L and q for which the following holds. Let $\xi \in L_q$, set H to be an L -subgaussian class and denote by $d_H = \sup_{h \in H} \|h\|_{L_2}$. For $w, u \geq 8$, with probability at least*

$$1 - c_0 w^{-q} N^{-((q/2)-1)} \log^q N - 2 \exp \left(-c_1 u^2 \left(\frac{\mathbb{E} \|G\|_H}{L d_H} \right)^2 \right),$$

$$\sup_{h \in H} \left| \frac{1}{N} \sum_{i=1}^N \xi_i h(X_i) - \mathbb{E} \xi h \right| \leq c_2 L w u \|\xi\|_{L_q} \frac{\mathbb{E} \|G\|_H}{\sqrt{N}}.$$

Proof of Lemma 3.1. The proof consists of two parts: first, controlling the process indexed by $\{f \in F : \|f - v_0\|_{L_2} \geq s\}$ where $s = (3/2)r_M(\eta_1, v_0)$, and then treating the process indexed by $\{v \in V : r \leq \|v - v_0\|_{L_2} \leq s\}$. Clearly, without loss of generality one may assume that $r \leq r_M(\eta_1, v_0)$.

By the regularity of r_M and since $s > r_M(\eta_1, v_0)$,

$$\mathbb{E} \|G\|_{(F-v_0) \cap sD} \leq \eta_1 \sqrt{N} s^2.$$

Moreover, $(F - v_0) \cap (s/4)D \subset (F - v_0) \cap sD$, and since $s/4 \leq r_M(\eta_1, v_0)$, the regularity of r_M implies that

$$\mathbb{E} \|G\|_{(F-v_0) \cap sD} \geq \eta_1 \sqrt{N} s^2 / 16.$$

Therefore, applying Theorem 3.2 to the set $H = (F - v_0) \cap sD$, there are constants c_1, c_2 and c_3 that depend only on q and L for which, with probability at least

$$1 - c_1 N^{-((q/2)-1)} \log^q N - 2 \exp \left(-c_2 \eta_1^2 s^2 N \right),$$

if $f \in F$ and $\|f - v_0\|_{L_2} \leq s$,

$$\left| \frac{1}{N} \sum_{i=1}^N \xi_i (f - v_0)(X_i) - \mathbb{E} \xi (f - v_0) \right| \leq c_3 L \|\xi\|_{L_q} \eta_1 s^2 = (*).$$

Clearly, $(*) \leq \theta s^2$ if $\eta_1 \leq \theta/c_3 L \|\xi\|_{L_q}$, and for such a choice, if $\|f - v_0\|_{L_2} = s$ then

$$\left| \frac{1}{N} \sum_{i=1}^N \xi_i(f - v_0)(X_i) - \mathbb{E}\xi(f - v_0) \right| \leq \theta \|f - v_0\|_{L_2}^2; \quad (3.4)$$

since $F - v_0$ is star-shaped around 0, (3.4) holds on the same event for every $f \in F$ for which $\|f - v_0\| \geq s$.

Next, one has to control the process indexed by $\{v \in V : r \leq \|v - v_0\|_{L_2} < s\}$. Set $j_0 = \lceil s/r \rceil$, fix $s_j = 2^j r$ for $0 \leq j \leq j_0$ and let $V_j = \text{star}((V - v_0) \cap s_j D)$. By Theorem 3.2, on an event \mathcal{A}_j , for every $h \in V_j$,

$$\left| \frac{1}{N} \sum_{i=1}^N \xi_i h(X_i) - \mathbb{E}\xi h \right| \leq c_4(L, q) w_j u_j \|\xi\|_{L_q} \frac{\mathbb{E}\|G\|_{V_j}}{\sqrt{N}} = (**)_j.$$

The aim is to ensure that $(**)_j \leq \theta s_j^2/4$ and that \mathcal{A}_j is of high enough probability. Indeed, on \mathcal{A}_j , if $v \in V$ and $s_j/2 \leq \|v - v_0\|_{L_2} \leq s_j$,

$$\left| \frac{1}{N} \sum_{i=1}^N \xi_i h(X_i) - \mathbb{E}\xi h \right| \leq \theta \|v - v_0\|_{L_2}^2.$$

To that end, let $w_j = \sqrt{j}$, recall that $d_V = \sup_{v \in V} \|v\|_{L_2}$ and thus $d_{V_j} = s_j = r 2^j$. Put

$$u_j = \max \left\{ 8, \frac{\sqrt{N}\theta}{4c_4\|\xi\|_{L_q}} \cdot \frac{2^j r}{\sqrt{j}} \cdot \frac{d_{V_j}}{\mathbb{E}\|G\|_{V_j}} \right\}$$

and consider two cases: first, if $u_j > 8$ then clearly, $(*) \leq \theta s_j^2/4$ and

$$Pr(\mathcal{A}_j) \geq 1 - c_5 \frac{\log^q N}{j^{q/2} N^{(q/2)-1}} - 2 \exp \left(-c_6(q, L) N \frac{2^{2j} \theta^2}{j \|\xi\|_{L_q}^2} \cdot r^2 \right).$$

Alternatively, if $u_j = 8$, then

$$u_j^2 \left(\frac{\mathbb{E}\|G\|_{V_j}}{d_{V_j}} \right)^2 \geq c_7(q, L) r^2 N \frac{2^{2j} \theta^2}{j \|\xi\|_{L_2}^2}.$$

Also, by (3.2), V_j has at most $|(V - v_0) \cap s_j D|$ extreme points. Since

$$\begin{aligned} \log |(V - v_0) \cap s_j D| &\leq c_8 \log(s_j/r) \log \mathcal{M}(F_{v_0} \cap 4rD, (r/2)D) \\ &\leq c_8 \log(s_j/r) \eta_1 \sqrt{N} r, \end{aligned}$$

by standard properties of gaussian processes

$$\begin{aligned}\mathbb{E}\|G\|_{V_j} &\leq c_9 d_{V_j} \cdot \log^{1/2} |(V - v_0) \cap s_j D| \leq c_{10} s_j \log^{1/2} \left(\frac{2s_j}{r} \right) \eta_1 \sqrt{Nr} \\ &= c_{10} \eta_1 \sqrt{N} \sqrt{\frac{j}{2^j}} s_j^2.\end{aligned}$$

Hence, there are constants c_{11} and c_{12} that depend only on q and L for which

$$\sup_{h \in V_j} \left| \frac{1}{N} \sum_{i=1}^N \xi_i h(X_i) - \mathbb{E} \xi h \right| \leq c_{11} \frac{\|\xi\|_{L_q}}{\sqrt{j}} \cdot \eta_1 \sqrt{\frac{j}{2^j}} s_j^2 \leq \theta s_j^2 / 4$$

if $\eta_1 \leq c_{12} \theta / \|\xi\|_{L_q}$.

Therefore, in both cases, there are constants c_{13} and c_{14} that depend only on q and L , and with probability at least

$$1 - c_{13} \frac{\log^q N}{j^{q/2} N^{(q/2)-1}} - 2 \exp(-c_{14} N r^2 \eta_1^2 2^j),$$

$$\sup_{h \in V_j} \left| \frac{1}{N} \sum_{i=1}^N \xi_i h(X_i) - \mathbb{E} \xi h \right| \leq \theta s_j^2 / 4.$$

The claim follows by applying the union bound to this estimate for $0 \leq j \leq j_0$. \blacksquare

Next, let us turn to the infimum of the quadratic process

$$\inf_{\{v \in V : \|v - v_0\|_{L_2} \geq cr\}} \frac{1}{N} \sum_{i=1}^N \left(\frac{(v - v_0)}{\|v - v_0\|_{L_2}} \right)^2 (X_i) \quad (3.5)$$

where r was selected in (3.3) for a well-chosen η_2 and where c is a suitable constant.

Lemma 3.3 *For every $L > 1$ there exist constants c_0, c_1 and c_2 that depend only on L for which the following holds. For every $v_0 \in V$, with probability at least $1 - 2 \exp(-c_0 N)$, if $v \in V$ and $\|v - v_0\|_{L_2} \geq c_1 r$ then*

$$\frac{1}{N} \sum_{i=1}^N (v - v_0)^2(X_i) \geq c_2 \|v - v_0\|_{L_2}^2.$$

The proof of Lemma 3.3 is similar to the one used in the analysis of the multiplier process: controlling relatively ‘large distances’ in F , i.e., when $f \in F$ for which $\|f - v_0\|_{L_2} \geq (3/2)r_Q(\eta_2) \equiv s$; and then ‘small distances’ in V , that is, $v \in V$ for which $r \leq \|v - v_0\|_{L_2} \leq s$ (again, one may assume that $r < r_Q(\eta_2)$).

For the constant η_2 (yet to be specified), one has

- $\mathbb{E}\|G\|_{(F-v_0) \cap sD} \leq \eta_2 N s$,
- for every $2r < t < s$,

$$\log \mathcal{N}((\text{star}(V - v_0)) \cap sD, tD) \leq c_0 \log(2s/t) \cdot \eta_2^2 N,$$

and

$$\log |(\text{star}(V - v_0)) \cap sD| \leq c_0 \log(2s/r) \cdot \eta_2^2 N.$$

The required lower bound on the infimum of the quadratic process (3.5) is based on estimates from [10] and [11], which will be formulated under the subgaussian assumption, rather than using the original (and much weaker) small-ball condition.

Theorem 3.4 *For every $L > 1$ there are constants κ_4 , κ_5 and κ_6 that depend only on L for which the following holds. Let H be an L -subgaussian class that is star-shaped around zero. Set $H_\rho = H \cap \rho D$ and fix ρ for which*

$$\mathbb{E}\|G\|_{H_\rho} \leq \kappa_4 \sqrt{N} \rho.$$

Then, with probability at least $1 - 2 \exp(-\kappa_5 N)$,

$$\inf_{\{h \in H : \|h\|_{L_2} \geq \rho\}} \frac{1}{N} \sum_{i=1}^N \left(\frac{h(X_i)}{\|h\|_{L_2}} \right)^2 \geq \kappa_6.$$

We will apply Theorem 3.4 to the class $H = (F - v_0) \cap sD$ (large distances) and then to $V_j = \text{star}((V - v_0) \cap s_j D)$ for $s_j = 2^j r$ (small distances).

Lemma 3.5 *There exist absolute constants c_0 and c_1 for which the following holds. For every $s > \rho \geq c_0 r$,*

$$\mathbb{E}\|G\|_{V_j \cap \rho D} \leq c_1 \eta_2 \sqrt{N} \left(\rho \log^{3/2}(2s_j/\rho) + r \log^{3/2}(2s/r) \right).$$

In particular, setting $\rho = s_j/2$ for $\eta_2 = c_2 \kappa_4$, one has

$$\mathbb{E}\|G\|_{V_j \cap (s_j/2)D} \leq \kappa_4 \sqrt{N} (s_j/2).$$

Proof. Fix $\rho < s_j$ and note that by Dudley's entropy integral bound (see, e.g., [8, 18]),

$$\begin{aligned} \mathbb{E}\|G\|_{V_j \cap \rho D} &\leq c_1 \int_0^\rho \log^{1/2} \mathcal{N}(V_j \cap \rho D, tD) dt \\ &\leq c_1 \int_0^r \log^{1/2} \mathcal{N}(V_j \cap \rho D, tD) dt + c_1 \int_r^\rho \log^{1/2} \mathcal{N}(V_j \cap \rho D, tD) dt. \end{aligned}$$

Applying (3.1) and since

$$V_j = \text{star}((V - v_0) \cap s_j D) \subset (\text{star}(V - v_0)) \cap s_j D,$$

it follows that for $r < t < \rho$,

$$\begin{aligned} \log \mathcal{N}(V_j \cap \rho D, rD) &\leq \log \mathcal{N}((\text{star}(V - v_0)) \cap \rho D, rD) \\ &\leq c_2 \log(2\rho/r) \cdot \sup_{x \in F} \log \mathcal{N}((F - x) \cap 4rD, rD) \\ &\leq c_2 \log(2\rho/r) \cdot \eta_2^2 N. \end{aligned}$$

Moreover, by (3.2),

$$\log |(V - v_0) \cap s_j D| \leq c_2 \log(2s_j/r) \cdot \eta_2^2 N = (*).$$

Hence, V_j is the union of at most $\exp(*)$ ‘intervals’ of the form $[0, v - v_0]$, and for $t \leq r$,

$$\log \mathcal{N}(V_j \cap \rho D, tD) \leq c_2 (\eta_2^2 N \log(2s_j/r) + \log(2\rho/t)).$$

Now the first part of the claim follows from integration, and the second part is an immediate outcome of the first. \blacksquare

Proof of Lemma 3.3. Combining Theorem 3.4 and Lemma 3.5 for $\eta_2 = c_0 \kappa_4$, it follows that with probability at least $1 - 2 \exp(-\kappa_5 N)$, if $v \in V$ and $s_j/2 \leq \|v - v_0\|_{L_2} \leq s_j$,

$$\frac{1}{N} \sum_{i=1}^N (v - v_0)^2(X_i) \geq \kappa_6 \|v - v_0\|_{L_2}^2. \quad (3.6)$$

Repeating this argument for $s_j = 2^j r$ and then applying it to the set $F_{v_0} \cap sD$ for $s = (3/2)r_Q(\eta_2)$, it follows that if $\log_2(s/r) \leq \exp(\kappa_5 N/2)$ then with probability at least $1 - 2 \exp(-\kappa_5 N/2)$, (3.6) holds for every $v \in V$ that satisfies $\|v - v_0\|_{L_2} \geq c_1 r$. \blacksquare

With all the ingredients in place, we may now conclude the proof of the upper estimate.

Fix $f_0 \in F$ and set $Y = f_0(X) + W$ for $W \in L_q$ that is orthogonal to $\text{span}(F)$. Let r , V and v_0 as above. Clearly, for every $v \in V$,

$$\|v - Y\|_{L_2}^2 = \|W\|_{L_2}^2 + \|v - f_0\|_{L_2}^2, \quad (3.7)$$

and thus $\|v_0 - f_0\|_{L_2} \leq r$. Moreover, for every $v \in V$, $\mathbb{E}W \cdot (v - v_0)(X) = 0$ and

$$\begin{aligned} |\mathbb{E}(v_0(X) - Y)(v - v_0)(X)| &= |\mathbb{E}(v_0 - f_0)(X) \cdot (v - v_0)(X)| \\ &\leq \|v_0 - f_0\|_{L_2} \cdot \|v - v_0\|_{L_2} \leq r\|v - v_0\|_{L_2}. \end{aligned}$$

By Lemma 3.3, with probability at least $1 - 2\exp(-\kappa_5 N/2)$, if $v \in V$ and $\|v - v_0\|_{L_2} \geq c(L)r$, then

$$\frac{1}{N} \sum_{i=1}^N (v - v_0)^2(X_i) \geq \kappa_6 \|v - v_0\|_{L_2}^2.$$

Using the notation of Lemma 3.1, set $\theta = \kappa_6/4$ and $\eta_1 = c_0(q, L)\theta/\|W\|_{L_q}$. Hence, there are constants c_1 and c_2 that depend only on q and L , for which, with probability at least

$$1 - c_1 \frac{\log^q N}{N^{((q/2)-1)}} - 2\exp(-c_2 \eta_1^2 r^2 N),$$

for every $v \in V$, $\|v - v_0\|_{L_2} \geq 2r$,

$$\left| \frac{1}{N} \sum_{i=1}^N (v_0(X_i) - Y_i)(v - v_0)(X_i) - \mathbb{E}(v_0(X) - Y)(v - v_0)(X) \right| \leq \frac{\kappa_6}{4} \|v - v_0\|_{L_2}^2.$$

On the intersection of the two events and for a constant $c_3 = c_3(q, L)$, if $\|v - v_0\|_{L_2} \geq c_3 r$ then

$$\begin{aligned} P_N \mathcal{L}_v^V &= \frac{1}{N} \sum_{i=1}^N (v - v_0)^2(X_i) + \frac{2}{N} \sum_{i=1}^N (v_0(X_i) - Y_i)(v - v_0)(X_i) \\ &\geq \frac{1}{N} \sum_{i=1}^N (v - v_0)^2(X_i) - 2|\mathbb{E}(v_0(X) - Y)(v - v_0)(X)| \\ &\quad - 2 \left| \frac{1}{N} \sum_{i=1}^N (v_0(X_i) - Y_i)(v - v_0)(X_i) - \mathbb{E}(v_0(X) - Y)(v - v_0)(X) \right| \\ &\geq \kappa_6 \|v - v_0\|_{L_2}^2 - 2r\|v - v_0\|_{L_2} - (\kappa_6/4)\|v - v_0\|_{L_2}^2 \geq (\kappa_6/4)\|v - v_0\|_{L_2}^2. \end{aligned}$$

Thus, for every such sample, the empirical minimizer $\hat{v} \in V$ satisfies that

$$\|\hat{v} - v_0\|_{L_2} \leq c_4 r.$$

And, since W is orthogonal to $\text{span}(F)$,

$$\begin{aligned} \mathbb{E}(\mathcal{L}_{\hat{v}}^F | (X_i, Y_i)_{i=1}^N) &= \|\hat{v} - Y\|_{L_2}^2 - \|f_0 - Y\|_{L_2}^2 = \|\hat{v} - f_0 - W\|_{L_2}^2 - \|W\|_{L_2}^2 \\ &= \|\hat{v} - f_0\|_{L_2}^2 - 2\mathbb{E}W \cdot (\hat{v} - f_0)(X) \leq (\|\hat{v} - v_0\|_{L_2} + \|v_0 - f_0\|_{L_2})^2 \leq (1 + c_4)^2 r^2. \end{aligned}$$

■

4 The lower bound

The lower estimates presented below are based on a volumetric argument. The idea is that if a learning procedure is ‘too successful’, a well-separated subset of F endows a well-separated subset in \mathbb{R}^N (a set that depends on X_1, \dots, X_N). However, because of some volumetric constraint, there is not ‘enough room’ for such a separated set to exist, leading to a contradiction.

The notions of volume are different in the two estimates: one is based on the Lebesgue measure while the other is determined by the choice of the ‘noise’ W , which is, in our case, gaussian.

Definition 4.1 *Let F be a class of functions and assume that $\mathbb{X} = (x_1, \dots, x_N) \in \Omega^N$. For every $f \in F$, set*

$$\mathcal{K}(f, \mathbb{X}) = \{h \in F : h(x_i) = f(x_i) \text{ for every } 1 \leq i \leq N\}.$$

The set $\mathcal{K}(f, \mathbb{X})$ is called the version space of F associated with f and \mathbb{X} .

In other words, $\mathcal{K}(f, \mathbb{X})$ consists of all the functions in F that agree with f on \mathbb{X} . Naturally, in the context of learning, \mathbb{X} is a random sample $(X_i)_{i=1}^N$, selected according to the underlying measure μ .

The diameter of the version space is a reasonable choice for a lower bound on the performance of any learning procedure: if $Y_i = f(X_i) + W_i$, a learning procedure cannot distinguish between f and any other function in the version space associated with f and $(X_i)_{i=1}^N$. Hence, the largest typical diameter of a version space should be a lower estimate on the performance of any learning procedure, as the following well-known fact shows (see, e.g., [7]).

Theorem 4.2 *Given a random variable W , for every $f \in F$ set $Y^f = f(X) + W$. If Ψ is a learning procedure, then*

$$\sup_{f \in F} \Pr \left(\|\Psi((Y_i^f, X_i)_{i=1}^N) - f\|_{L_2(\mu)} \geq \frac{1}{4} \mathcal{K}(f, \mathbb{X}) \right) \geq 1/2,$$

where the probability is relative to the product measure endowed on $(\Omega \times \mathbb{R})^N$ by the N -product of the joint distribution of X and W .

Clearly, if W is orthogonal to $\text{span}(F)$, then for every $h \in F$ and every target Y^f , $\mathbb{E} \mathcal{L}_h^f = \|h - f\|_{L_2}^2$. Thus, the largest typical diameter of a version space $\mathcal{K}(f, \mathbb{X})$ is a lower bound on \mathcal{E}_p for the set of admissible targets $\mathcal{Y} = \{f(X) + W : f \in F\}$.

This leads to the following question:

Question 4.3 *Given a class F defined on a probability space (Ω, μ) , $f \in F$ and $\mathbb{X} = (x_1, \dots, x_N) \subset \Omega^N$, find a lower estimate on*

$$\text{diam}(\mathcal{K}(f, \mathbb{X}), L_2(\mu)).$$

One situation in which Question 4.3 is of independent interest is when $T \subset \mathbb{R}^n$ is a convex body (i.e., a convex, centrally-symmetric set with a nonempty interior) and $F = \{\langle t, \cdot \rangle : t \in T\}$ is the class of linear functionals associated with T . For every $x_1, \dots, x_N \in \mathbb{R}^n$ set $\mathbb{X} = (x_1, \dots, x_N)$, and let $\Gamma_{\mathbb{X}} = \sum_{i=1}^N \langle x_i, \cdot \rangle e_i$ be the matrix whose rows are x_1, \dots, x_N . Thus,

$$\mathcal{K}(0, \mathbb{X}) = \ker(\Gamma_{\mathbb{X}}) \cap T.$$

If μ is an isotropic, L -subgaussian measure on \mathbb{R}^n , one may show that with probability at least $1 - 2 \exp(-c_0 N)$,

$$\text{diam}(\mathcal{K}(0, \mathbb{X}), L_2(\mu)) \leq 2r_Q(c_1(L)) \quad (4.1)$$

(see [14]). This extends the celebrated result of Pajor and Tomczak-Jaegermann [15, 16], that (4.1) holds for the Haar measure on S^{n-1} (and thus, also for the gaussian measure on \mathbb{R}^n).

It turns out that (4.1) is not far from optimal:

Theorem 4.4 *There exists an absolute constant c for which the following holds. Let $F \subset L_2(\mu)$ be a convex and centrally-symmetric set. If*

$$\log \mathcal{M}(F \cap 2rD, (r/4)D) \geq cN,$$

then for every $\mathbb{X} = (x_1, \dots, x_N)$,

$$\text{diam}(\mathcal{K}(0, \mathbb{X}), L_2(\mu)) \geq r/8.$$

Since F is convex and centrally-symmetric, $F - F = 2F$ and $0 \in F$. Therefore,

$$\begin{aligned} \mathcal{M}(F \cap 4rD, (r/2)D) &\leq \sup_{x \in F} \mathcal{M}((F - x) \cap 4rD, (r/2)D) \\ &\leq \mathcal{M}((F - F) \cap 4rD, (r/2)D) = \mathcal{M}(F \cap 2rD, (r/4)D). \end{aligned}$$

Hence, Theorem 4.4 shows that if $\gamma_Q(c, 0) > r$ then for every $\mathbb{X} = (x_1, \dots, x_N)$, $\text{diam}(\mathcal{K}(0, \mathbb{X}), L_2(\mu)) \geq r/8$. In particular, for every $W \in L_2$ that is orthogonal to $\text{span}(F)$, the best possible error rate in F that holds for every target $Y^f = f(X) + W$, is at least $\gamma_Q^2(c, 0) \geq c_1 \gamma_Q^2(c)$.

Proof. Let f_1, \dots, f_m be $r/4$ -separated in $F \cap 2rD$. Set

$$A_i = \frac{f_i}{2} + \frac{1}{32}(F \cap 2rD),$$

and observe that $A_i \subset F \cap 2rD$. Also, for every $h \in A_i$, $\|(f_i/2) - h\|_{L_2} \leq r/16$; therefore, if $h_i \in A_i$ and $h_\ell \in A_\ell$, then $\|h_i - h_\ell\|_{L_2} \geq r/8$.

Fix $\mathbb{X} = (x_1, \dots, x_N)$ and for $A \subset F$ set

$$P_{\mathbb{X}}(A) = \{(h(X_i))_{i=1}^N : h \in A\} \subset \mathbb{R}^N,$$

the coordinate projection of A associated with \mathbb{X} . Clearly, for every $1 \leq i \leq m$,

$$P_{\mathbb{X}}(A_i) = \frac{1}{2}(f_i(x_j))_{j=1}^N + \frac{1}{32}P_{\mathbb{X}}(F \cap 2rD). \quad (4.2)$$

Consider two possibilities. First, if there are $i \neq \ell$ for which $P_{\mathbb{X}}(A_i) \cap P_{\mathbb{X}}(A_\ell) \neq \emptyset$, there are $h_i \in A_i$ and $h_\ell \in A_\ell$ that satisfy $h_i - h_\ell \in \mathcal{K}(0, \mathbb{X})$, thus showing that $\text{diam}(\mathcal{K}(0, \mathbb{X}), L_2(\mu)) \geq r/8$.

Otherwise, the sets $P_{\mathbb{X}}(A_i)$ are disjoint subsets of $P_{\mathbb{X}}(F \cap 2rD)$. And, setting $T = P_{\mathbb{X}}(F \cap 2rD)$, (4.2) implies that $\mathcal{M}(T, T/32) \geq m$. Since T is a convex, centrally symmetric subset of \mathbb{R}^N , a standard volumetric argument shows that $\mathcal{M}(T, T/32) \leq \exp(cN)$ for a suitable absolute constant c . Thus, if $m > \exp(cN)$, $\text{diam}(\mathcal{K}(0, \mathbb{X}), L_2) \geq r/8$, as claimed. \blacksquare

The final result of this section is the ‘noise-dependent’ lower bound.

Theorem 4.5 *There exist absolute constants c_1 and c_2 for which the following holds. Let $F \subset L_2(\mu)$ be a convex, centrally-symmetric class of functions, set W to be a centred normal random variable that is independent of X , and for every $f \in F$, put $Y^f = f(X) + W$. If Ψ is a learning procedure that*

performs with confidence of at least $7/8$ for every Y^f , there is some Y^f for which

$$\mathcal{E}_p \geq c_1 \gamma_M^2 \left(\frac{c_2}{\|W\|_{L_2}} \right).$$

Stronger versions of Theorem 4.5 (without the assumption that F is convex and centrally-symmetric) may be proved in several different ways: using information theoretic tools (see, Theorem 2.5 in [17]), or, alternatively, by applying the gaussian isoperimetric inequality as in [7]. Both these arguments are rather restrictive, because they rely on rather special properties of the noise.

Although the proof we present below is also for a gaussian noise, the argument is less restrictive and may be extended to other choices of noise (e.g. when W is log-concave rather than gaussian). The argument is essentially the same as Talagrand's proof of the dual-Sudakov inequality [8], and as such is volumetric in nature: obtaining a lower bound on the measure of a shift of a centrally-symmetric set in terms of the Euclidean norm of the shift.

Lemma 4.6 *Let $A \subset \mathbb{R}^N$ be centrally symmetric and set $z \in \mathbb{R}^N$. If ν is the centred gaussian measure on \mathbb{R}^N with covariance $\sigma^2 I_N$ and $|\cdot|$ denotes the Euclidean norm on \mathbb{R}^n , then*

$$\nu(z + A) \geq \exp \left(-\frac{|z|^2}{2\sigma^2} \right) \nu(A).$$

Proof. A change of variables shows that

$$\begin{aligned} \nu(z + A) &= \frac{1}{(2\pi\sigma)^{N/2}} \int_{z+A} \exp \left(-\frac{|x|^2}{2\sigma^2} \right) dx = \frac{1}{(2\pi\sigma)^{N/2}} \int_A \exp \left(-\frac{|t+z|^2}{2\sigma^2} \right) dt \\ &= \exp \left(-\frac{|z|^2}{2\sigma^2} \right) \cdot \frac{1}{(2\pi\sigma)^{N/2}} \int_A \exp \left(\frac{\langle z, t \rangle}{\sigma^2} \right) \cdot \exp \left(-\frac{|t|^2}{2\sigma^2} \right) dt = (*). \end{aligned}$$

Let $\mathbb{E}_{\nu|A}$ be the expectation with respect to the gaussian measure ν , conditioned on A . Thus,

$$(*) = \exp \left(-\frac{|z|^2}{2\sigma^2} \right) \nu(A) \cdot \mathbb{E}_{\nu|A} \exp \left(-\frac{\langle z, t \rangle}{\sigma^2} \right).$$

Since A is symmetric, $\mathbb{E}_{\nu|A} \langle z, t \rangle = 0$, and by Jensen's inequality

$$(*) \geq \exp \left(-\frac{|z|^2}{2\sigma^2} \right) \nu(A).$$

■

Proof of Theorem 4.5. Let Ψ be a learning procedure that performs with accuracy \mathcal{E}_p for every target $Y^f = f(X) + W$ for $f \in F$ and $W \sim \mathcal{N}(0, \sigma^2)$ that is independent of X . Note that for the target Y^f , the true minimizer in F is $f^* = f$ and for every $h \in F$,

$$\mathbb{E}\mathcal{L}_h = \mathbb{E}(h(X) - Y^f)^2 - \mathbb{E}(f^*(X) - Y^f)^2 = \|h - f\|_{L_2}^2.$$

Thus, if $\tau = (x_i, y_i)_{i=1}^N \in (\Omega \times \mathbb{R})^N$ is a sample on which Ψ performs with accuracy \mathcal{E}_p relative to the target Y^f , then $\|\Psi(\tau) - f^*\|_{L_2}^2 \leq \mathcal{E}_p$.

Let $(f_j)_{j=1}^m$ be a subset of $F \cap 4rD$ that is $r/2$ separated in $L_2(\mu)$ for $(r/2)^2 = 9\mathcal{E}_p$ and fix $\mathbb{X} = (x_1, \dots, x_N) \in \Omega^N$.

For every $1 \leq j \leq m$, put

$$A_j(\mathbb{X}) = \left\{ (w_i)_{i=1}^N : \Psi((x_i, f_j(x_i) + w_i)_{i=1}^N) \in f_j + \sqrt{\mathcal{E}_p}D \right\} \subset \mathbb{R}^N,$$

i.e., $A_j(\mathbb{X})$ consists of all the vectors $(w_i)_{i=1}^N \in \mathbb{R}^N$, for which, upon receiving the data $(x_i, f_j(x_i) + w_i)_{i=1}^N$, Ψ selects a point whose L_2 distance to f_j is at most $r/6 = \sqrt{\mathcal{E}_p}$.

Let ν be the centred gaussian measure on \mathbb{R}^N with covariance $\sigma^2 I_N$. Since W is a centred gaussian random variable with variance σ^2 , $(w_i)_{i=1}^N$ is distributed according to ν , and since it is independent of X , if Ψ performs with accuracy \mathcal{E}_p and with probability at least $7/8$, it is evident that

$$\begin{aligned} & \mu^N \otimes \nu \left(\left\{ (x_i, w_i)_{i=1}^N : \Psi((x_i, f_j(x_i) + w_i)_{i=1}^N) \in f_j + \sqrt{\mathcal{E}_p}D \right\} \right) \\ &= \mu^N \otimes \nu \left(\left\{ (x_i, w_i)_{i=1}^N : (w_i)_{i=1}^N \in A_j(\mathbb{X}) \right\} \right) \geq 7/8. \end{aligned}$$

A standard Fubini argument shows that there is an event $\mathcal{C}_j \subset \Omega^N$ of μ^N probability at least $1/2$, and for every $\mathbb{X} = (x_i)_{i=1}^N \in \mathcal{C}_j$, $\nu(A_j(\mathbb{X})) \geq 3/4$. Observe that if $\mathbb{X} \in \mathcal{C}_j$ then by the symmetry of ν , $\nu(-A_j(\mathbb{X})) \geq 3/4$, and the centrally-symmetric set $A_j(\mathbb{X}) \cap -A_j(\mathbb{X}) \subset A_j(\mathbb{X})$ satisfies that

$$\nu(A_j(\mathbb{X}) \cap -A_j(\mathbb{X})) \geq 1/2.$$

Let $z_j = (f_j(x_i))_{i=1}^N$. If $\mathbb{X} \in \mathcal{C}_j \cap \mathcal{C}_\ell$, the sets $z_j + A_j(\mathbb{X})$ and $z_\ell + A_\ell(\mathbb{X})$ are disjoint, because Ψ maps $z_j + A_j(\mathbb{X})$ to an $r/6$ -neighbourhood of f_j and $z_\ell + A_\ell(\mathbb{X})$ to an $r/6$ -neighbourhood of f_ℓ – but $\|f_j - f_\ell\|_{L_2} \geq r/2$. Therefore

$$\sum_{j=1}^m \mathbb{1}_{\mathcal{C}_j}(\mathbb{X}) \nu(z_j + (A_j(\mathbb{X}) \cap -A_j(\mathbb{X}))) \leq 1;$$

integrating with respect to μ^N ,

$$\sum_{i=1}^m \mathbb{E}_{\mathbb{X}} \mathbb{1}_{\mathcal{C}_j}(\mathbb{X}) \nu(z_j + (A_j(\mathbb{X}) \cap -A_j(\mathbb{X}))) \leq 1,$$

and all that remains is to control $\mathbb{E}_{\mathbb{X}} \mathbb{1}_{\mathcal{C}_j}(\mathbb{X}) \nu(z_j + (A_j(\mathbb{X}) \cap -A_j(\mathbb{X})))$ from below.

Applying Lemma 4.6,

$$\begin{aligned} \nu(z_j + (A_j(\mathbb{X}) \cap -A_j(\mathbb{X}))) &\geq \exp\left(-\frac{|z_j|^2}{2\sigma^2}\right) \cdot \nu(A_j(\mathbb{X}) \cap -A_j(\mathbb{X})) \\ &= \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^N f_j^2(x_i)\right) \cdot \nu(A_j(\mathbb{X}) \cap -A_j(\mathbb{X})). \end{aligned}$$

By Chebychev's inequality and recalling that $\|f_j\|_{L_2} \leq 4r$,

$$\mu^N \left\{ \sum_{i=1}^N f_j^2(X_i) \leq c_0 N r^2 \right\} \geq 3/4$$

for an appropriate choice of an absolute constant c_0 and for every $1 \leq j \leq m$. Thus, on an event of μ^N measure at least $1/4$, $\mathbb{X} \in \mathcal{C}_j$, $\nu(A_j(\mathbb{X}) \cap -A_j(\mathbb{X})) \geq 3/4$ and $\sum_{i=1}^N f_j^2(X_i) \leq c_0 N r^2$; therefore,

$$\mathbb{E}_{\mathbb{X}} \mathbb{1}_{\mathcal{C}_j}(\mathbb{X}) \nu(z_j + (A_j(\mathbb{X}) \cap -A_j(\mathbb{X}))) \geq c_1 \exp\left(-c_0 \frac{N r^2}{2\sigma^2}\right).$$

Hence, $\log m \leq c_2^2 N r^2 / \sigma^2$, i.e., $\log \mathcal{M}(F \cap 4rD, (r/2)D) \leq (c_2/\sigma)^2 N r^2$, implying that $\mathcal{E}_p \geq c_3 \gamma_M^2 (c_2/\sigma)$. \blacksquare

5 Some Remarks

We begin this section with an example of ‘natural’ sets, for which there is a true gap between the two sets of parameters: r_Q/r_M and γ_Q/γ_M .

Let $T \subset \mathbb{R}^n$ be a convex body in \mathbb{R}^n (i.e., a convex, centrally-symmetric set with a nonempty interior), put $F = \{\langle t, \cdot \rangle : t \in T\}$, the class of linear functionals associated with T and set μ to be the gaussian measure on \mathbb{R}^n .

It is straightforward to verify that for every $r > 0$, $(F \cap rD, L_2(\mu))$ is isometric to $(T \cap rB_2^n, \ell_2^n)$, where B_2^n is the Euclidean unit ball in \mathbb{R}^n . Let $1 \leq p < 2$, and set $T = B_p^n$, the unit ball in $\ell_p^n = (\mathbb{R}^n, \|\cdot\|_{\ell_p})$. One may

show (see [7]) that when $p = 1$, r_M and γ_M are equivalent, as are r_Q and γ_Q . However, such an equivalence is no longer true for $1 < p < 2$ (of course, as long as $p > 1 + 1/\log n$ – otherwise, ℓ_p^n is equivalent to ℓ_1^n).

To see how that gap between the ‘global’ and ‘local’ parameters is exhibited in B_p^n for $1 < p < 2$, let $x = (x_i)_{i=1}^n \in B_p^n$ and set $(x_i^*)_{i=1}^n$ to be the non-increasing rearrangement of $(|x_i|)_{i=1}^n$; thus, $x_i^* \leq i^{-1/p}$. Recall the well known fact (see, e.g., [5]), that $\mathbb{E}\|G\|_{B_p^n \cap rB_2^n}$ is equivalent to

$$c_1(p) \begin{cases} n^{1-1/p} & \text{if } r \geq c_2(p)n^{-(1/p-1/2)}, \\ rn^{1/2} & \text{if } r \leq c_2(p)n^{-(1/p-1/2)}. \end{cases}$$

Thus, if $N \leq n^{2/p}$,

$$r_M \sim n^{1/2-1/2p}/N^{1/4} \geq c_2(p)n^{-(1/p-1/2)}.$$

Let us consider the case in which $1 > r \gg c_2(p)n^{-(1/p-1/2)}$. Set $\ell = (1/r)^{2p/(2-p)}$ and observe that

$$B_p^n \cap 4rB_2^n \subset \left\{ x \in \mathbb{R}^n : x_i^* \leq 4r/i^{1/2} \text{ if } i \leq \ell, \text{ and } x_i^* \leq 1/i^{1/p} \text{ if } i > \ell \right\}.$$

Clearly, for a well-chosen constant c_3 one has $\sum_{i \geq c_3\ell} i^{-2/p} \leq r^2/100$, and

$$B_p^n \cap 4rB_2^n \subset \bigcup_{|I|=c_3\ell} (c_4rB_{2,\infty}^I + B_{p,\infty}^{I^c}),$$

where $B_{q,\infty}^I$ is the unit ball in \mathbb{R}^I endowed with the weak $\ell_{q,\infty}$ norm¹. In particular, if $|I| = c_3\ell$, $B_{p,\infty}^{I^c} \subset (r/10)B_2^{I^c}$ and the impact of those ‘small’ coordinates on Euclidean distances is negligible:

$$\mathcal{M}(B_p^n \cap 4rB_2^n, rB_2^n) \leq \mathcal{M}\left(\bigcup_{|I|=c_3\ell} c_4rB_{2,\infty}^I, (r/2)B_2^n\right) \leq \binom{n}{c_3\ell} c_4^{c_3\ell}.$$

Hence, separation at scale r occurs only because of the largest $\sim \ell$ coordinates of the vectors involved.

On the other hand, the contribution of those ‘large’ coordinates to $\mathbb{E}\|G\|_{B_p^n \cap 4rB_2^n}$ is equally negligible. Indeed, if $T = \bigcup_{|I|=m} \alpha rB_2^I$ for some $m \leq n/2$ and $\alpha \geq 1$, it is standard to verify that

$$\mathbb{E} \sup_{t \in T} \sum_{i=1}^n g_i t_i \leq c_5 \alpha r m^{1/2} \cdot \log^{1/2} \left(\frac{en}{m} \right) = (*).$$

¹Recall that for $x \in \mathbb{R}^n$, $\|x\|_{q,\infty} \leq A$ if and only if, $\sup_{i \geq 1} i^{1/q} x_i^* \leq A$.

If $m = c_3 \ell = c_3(1/r)^{2p/(2-p)}$ and since $r \gg c_2(p)n^{-(1/p-1/2)}$, it follows that

$$(*) \ll n^{1-1/p} \sim \mathbb{E}\|G\|_{B_p^n \cap rB_2^n}.$$

Thus, as long as r is significantly larger than $c_2(p)n^{-(1/p-1/2)}$, the gaussian average of the intersection body $B_p^n \cap 4rB_2^n$ originates from the ‘small coordinates’ in the monotone rearrangement, and in particular, from vectors whose Euclidean norm is significantly smaller than r . Such vectors are ‘invisible’ to γ_M , which is why γ_M is much smaller than r_M .

5.1 The role of fixed points

Fixed points are encountered frequently in Empirical Processes and Statistics literature, and almost always with the same goal: obtaining ‘relative’ upper bounds on various empirical processes. To obtain such bounds, one has to compare the oscillation (i.e., the behaviour of the process indexed by $(F - F) \cap rD$) with some function of r .

One usually obtains upper bounds on the oscillation via a symmetrization argument, leading to a sample-dependent Bernoulli process. Thus, the standard outcome is a fixed point equation, linking an entropy integral relative to the random L_2 metric and generated by the sample X_1, \dots, X_N , with the desired function of r (see [18] for numerous examples).

Still within the realm of entropy integrals, it is possible to impose additional structure on the problem, which allows one to replace the empirical L_2 (random) metrics with the global $L_2(\mu)$ metric. For example, a fixed point equation with the same normalization as r_M may be found in [2], where the setup allows the transition between the random metric and the deterministic one – but the ‘philosophy’ of the proof is the same: it is based on an entropy integral.

Since the entropy integral is only upper estimate on the supremum of the empirical process in question – regardless of the underlying assumptions, it is often loose. Therefore, one would like to find a general argument bypassing the whole mechanism of entropy integrals.

As a first step, and because it is natural to expect that the empirical processes in question converges to a gaussian limit, one may try a ‘gaussian’-based fixed point, which relies on $\mathbb{E}\|G\|_{(F-F) \cap rD}$, rather than on an entropy integral bound. And, indeed, under a subgaussian assumption, the results of [7] lead to the gaussian-based r_M and r_Q .

Our results show that r_M and r_Q are not the end of the story and can be improved – at least for the special learning problems we consider. The

‘right’ fixed points should involve the smaller local entropy estimates rather than the oscillation of the gaussian process.

One fixed point that seems closer in nature to γ_M than to r_M may be found in the celebrated work of Yang and Barron [19], though a closer inspection shows that this impression is inaccurate.

Comparing [19] to our results is somewhat unnatural because the setup in [19] is completely different: a function class consisting of uniformly bounded functions and an independent gaussian noise, both of which are crucial to the proof (see Section 3.2 in [19]). Also, the upper estimate is an existence result of a ‘good’ procedure – rather than a specific choice of a procedure; the estimate holds in expectation and not with high probability; and it does not tend to zero with the ‘noise level’ of the problem.

All these differences are significant, but are still not a conclusive indication that the nature of the complexity parameter in [19] is different from ours. That indication is the key to the results in [19]: the assumption that the underlying class ‘large’ – in the sense that

$$\liminf_{\varepsilon \rightarrow 0} \frac{\log \mathcal{M}(F, (\varepsilon/2)D)}{\log \mathcal{M}(F, \varepsilon D)} > 1. \quad (5.1)$$

One should note that this assumption immediately excludes all the modern high-dimensional problems, involving classes indexed by subsets of \mathbb{R}^n . Indeed, for any convex subset of \mathbb{R}^n , the liminf above is 1 rather than strictly greater than 1.

Equation (5.1) has two significant implications:

- The $r/2$ log-covering numbers of F and of $F \cap rD$ are equivalent, which means that one may replace the local sets $F \cap rD$ with F in the definition of the fixed points. This makes the proof of the upper bound simpler.
- It essentially restricts the setup to classes that have polynomial entropy, which is a considerably narrower scenario. Indeed, for the sake of brevity let us ignore cases in which

$$\limsup_{\varepsilon \rightarrow 0} \frac{\log \mathcal{M}(F, (\varepsilon/2)D)}{\log \mathcal{M}(F, \varepsilon D)} = L \geq 4$$

(if $L > 4$ then the gaussian process $\{G_f : f \in F\}$ is not bounded and the class F is not subgaussian, while if $L = 4$ an entropy estimate is not enough to determine whether the gaussian process is bounded and

thus requires a more subtle analysis). When $L < 4$ and because the entropy of the ‘local’ set $F \cap rD$ is equivalent to the entropy of F , it follows that there are $0 < q_1 \leq q_2 < 2$ for which, for every $\varepsilon \leq R$ small enough,

$$\left(\frac{R}{\varepsilon}\right)^{q_1} \lesssim \frac{\log \mathcal{M}(F \cap RD, (\varepsilon/2)D)}{\log \mathcal{M}(F \cap RD, \varepsilon D)} \lesssim \left(\frac{R}{\varepsilon}\right)^{q_2}.$$

Using Dudley’s entropy integral for the upper bound and Sudakov’s minoration for the lower one, it is straightforward to verify that

$$\mathbb{E}\|G\|_{F \cap RD} \sim_{q_1, q_2} R \log^{1/2} \mathcal{M}(F \cap RD, (R/2)D).$$

Therefore, the ‘global’ parameters r_M and r_Q are equivalent to the local ones γ_M and γ_Q ; in fact, the ‘local’ and ‘global’ parameters are even equivalent to the ones defined via the entropy integral. Thus, the typical situation in [19] is very different from the problems studied here – mainly because of (5.1).

References

- [1] Martin Anthony and Peter L. Bartlett. *Neural network learning: theoretical foundations*. Cambridge University Press, Cambridge, 1999.
- [2] Lucien Birgé and Pascal Massart. Rates of convergence for minimum contrast estimators. *Probab. Theory Related Fields*, 97(1-2):113–150, 1993.
- [3] Peter Bühlmann and Sara A. van de Geer. *Statistics for high-dimensional data*. Springer Series in Statistics. Springer, Heidelberg, 2011. Methods, theory and applications.
- [4] R. M. Dudley. *Uniform central limit theorems*, volume 63 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999.
- [5] Yehoram Gordon, Alexandre E. Litvak, Shahar Mendelson, and Alain Pajor. Gaussian averages of interpolated bodies and applications to approximate reconstruction. *J. Approx. Theory*, 149(1):59–73, 2007.
- [6] Vladimir Koltchinskii. *Oracle inequalities in empirical risk minimization and sparse recovery problems*, volume 2033 of *Lecture Notes in Mathematics*. Springer, Heidelberg, 2011. Lectures from the 38th Probability Summer School held in Saint-Flour, 2008, École d’Été de Probabilités de Saint-Flour. [Saint-Flour Probability Summer School].
- [7] Guillaume Lecué and Shahar Mendelson. Learning subgaussian classes: Upper and minimax bounds. Technical report, CNRS, Ecole polytechnique and Technion, 2013.
- [8] Michel Ledoux and Michel Talagrand. *Probability in Banach spaces*, volume 23 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1991. Isoperimetry and processes.

- [9] Pascal Massart. *Concentration inequalities and model selection*, volume 1896 of *Lecture Notes in Mathematics*. Springer, Berlin, 2007. Lectures from the 33rd Summer School on Probability Theory held in Saint-Flour, July 6–23, 2003, With a foreword by Jean Picard.
- [10] Shahar Mendelson. Learning without concentration. *Journal of the ACM*. To appear.
- [11] Shahar Mendelson. Learning without concentration for general loss functions. Arxiv: <http://arxiv.org/abs/1410.3192>.
- [12] Shahar Mendelson. Upper bounds on product and multiplier empirical processes. Arxiv: <http://arxiv.org/abs/1410.8003>.
- [13] Shahar Mendelson. Obtaining fast error rates in nonconvex situations. *J. Complexity*, 24(3):380–397, 2008.
- [14] Shahar Mendelson, Alain Pajor, and Nicole Tomczak-Jaegermann. Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geom. Funct. Anal.*, 17(4):1248–1282, 2007.
- [15] A. Pajor and N. Tomczak-Jaegermann. Nombres de Gel’fand et sections euclidiennes de grande dimension. In *Séminaire d’Analyse Fonctionnelle 1984/1985*, volume 26 of *Publ. Math. Univ. Paris VII*, pages 37–47. Univ. Paris VII, Paris, 1986.
- [16] Alain Pajor and Nicole Tomczak-Jaegermann. Subspaces of small codimension of finite-dimensional Banach spaces. *Proc. Amer. Math. Soc.*, 97(4):637–642, 1986.
- [17] Alexandre B. Tsybakov. *Introduction to nonparametric estimation*. Springer Series in Statistics. Springer, New York, 2009. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats.
- [18] Aad W. van der Vaart and Jon A. Wellner. *Weak convergence and empirical processes*. Springer Series in Statistics. Springer-Verlag, New York, 1996. With applications to statistics.
- [19] Yuhong Yang and Andrew Barron. Information-theoretic determination of minimax rates of convergence. *Ann. Statist.*, 27(5):1564–1599, 1999.